



Europäisches Patentamt  
European Patent Office  
Office européen des brevets



(11) EP 1 024 428 A2

(12)

## EUROPEAN PATENT APPLICATION

(43) Date of publication:  
02.08.2000 Bulletin 2000/31

(51) Int. Cl.<sup>7</sup>: G06F 9/50, G06F 11/00

(21) Application number: 00300598.0

(22) Date of filing: 27.01.2000

(84) Designated Contracting States:  
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU  
MC NL PT SE  
Designated Extension States:  
AL LT LV MK RO SI

• Goal, Patrick M.,  
IBM United Kingdom Ltd., I.P.L.  
Winchester, Hampshire SO21 2JN (GB)  
• McCarty, Richard James,  
U. K. Ltd., I.P.L.  
Winchester, Hampshire SO21 2JN (GB)

(30) Priority: 29.01.1999 US 240494

(71) Applicant:  
International Business Machines Corporation  
Armonk, NY 10504 (US)

(74) Representative:  
Davies, Simon Robert  
IBM  
UK Intellectual Property Department  
Hursley Park  
Winchester, Hampshire SO21 2JN (GB)

(72) Inventors:  
• Chao, Ching-Yun,  
IBM United Kingdom Ltd., I.P.L.  
Winchester, Hampshire SO21 2JN (GB)

### (54) Managing a clustered computer system

(57) A clustered computer system provides both speed and reliability advantages. However, when communications between the clustered computers are compromised those same computers can become confused and corrupt database files. The present method and apparatus are used to improve the management of clus-

tered computer systems. Specifically, the system expands the number of nodes available for failover conditions. Further, provision is made for returning the system to an initial state after a failover event.




FIG. 3

## Description

[0001] The present invention relates generally to a distributed data processing system and in particular to a method and apparatus for managing a clustered computer system.

[0002] A clustered computer system is a type of parallel or distributed system that consists of a collection of interconnected whole computers and is used as a single, unified computing resource. The term "whole computer" in the above definition is meant to indicate the normal combination of elements making up a stand-alone, usable computer: one or more processors, an acceptable amount of memory, input/output facilities, and an operating system. Another distinction between clusters and traditional distributed systems concerns the relationship between the parts. Modern distributed systems use an underlying communication layer that is peer-to-peer. There is no intrinsic hierarchy or other structure, just a flat list of communicating entities. At a higher level of abstraction, however, they are popularly organized into a client-server paradigm. This results in a valuable reduction in system complexity. Clusters typically have a peer-to-peer relationship.

[0003] There are three technical trends to explain the popularity of clustering. First, microprocessors are increasingly fast. The faster microprocessors become, the less important massively parallel systems become. It is no longer necessary to use super-computers or aggregations of thousands of microprocessors to achieve suitably fast results. A second trend that has increased the popularity of clustered computer systems is the increase in high-speed communications between computers. A cluster computer system is also referred to as a "cluster". The introduction of such standardized communication facilities as Fibre Channel Standard (FCS), Asynchronous Transmission Mode (ATM), the Scalable Coherent Interconnect (SCI), and the switched Gigabit Ethernet are raising inter-computer bandwidth from 10 Mbits/second through hundreds of Mbytes/second and even Gigabytes per second. Finally, standard tools have been developed for distributed computing. The requirements of distributed computing have produced a collection of software tools that can be adapted to managing clusters of machines. Some, such as the Internet communication protocol suite (called TCP/IP and UDP/IP) are so common as to be ubiquitous de facto standards. High level facilities built on the base, such as Intranets, the Internet and the World Wide Web, are similarly becoming ubiquitous. In addition, other tool sets for multi-sense administration have become common. Together, these are an effective base to tap into for creating cluster software.

[0004] In addition to these three technological trends, there is a growing market for computer clusters. In essence, the market is asking for highly reliable computing. Another way of stating this is that the computer networks must have "high availability". For example, if

the computer is used to host a web-site, its usage is not necessarily limited to normal business hours. In other words, the computer may be accessed around the clock, for every day of the year. There is no safe time to shut down to do repairs. Instead, a clustered computer system is useful because if one computer in the cluster shuts down, the others in the cluster automatically assume its responsibilities until it can be repaired. There is no down-time exhibited or detected by users.

[0005] Businesses need high availability for other reasons as well. For example, business-to-business intranet use involves connecting businesses to subcontractors or vendors. If the intranet's file servers go down, work by multiple companies is strongly affected. If a business has a mobile workforce, that workforce must be able to connect with the office to download information and messages. If the office's server goes down, the effectiveness of that work force is diminished.

[0006] A computer system is highly available when no replaceable piece is a single point of failure, and overall, it is sufficiently reliable that one can repair a broken part before something else breaks. The basic technique used in cluster to achieve high availability is failover. The concept is simple enough: one computer (A) watches over another computer (B); if B dies, A takes over B's work. Thus, failover involves moving "resources" from one node to another. A node is another term for a computer. Many different kinds of things are potentially involved: physical disk ownership, logical disk volumes, IP addresses, application processes, subsystems, print queues, collection of cluster-wide locks in a shared-data system, and so on.

[0007] Resources depend on one another. The relationship matters because, for example, it will not help to move an application to one node when the data it uses is moved to another. Actually it will not even help to move them both to the same node if the application is started before the necessary disk volumes are mounted.

[0008] In modern cluster systems such as IBM HACMP and Microsoft "Wolfpack", the resource relationship information is maintained in a cluster-wide data file. Resources that depend upon one another are organized as a resource group and are stored as a hierarchy in that data file. A resource group is the basic unit of failover.

[0009] With reference now to the figures, and in particular with reference to Figure 1, this provides a pictorial representation of a distributed data processing system 100 including a network of computers. Distributed data processing system 100 contains one or more public networks 101, which is the medium used to provide communications links between various devices, client computers, and server computers connected within distributed data processing system 100. Network 100 may include permanent connections, such as Token Ring, Ethernet, 100Mb Ethernet, Gigabit Ethernet, FDDI ring, ATM, and high speed switch, or temporary

connections made through telephone connections. Client computers 130 and 131 communicate to server computers 110, 111, 112, and 113 via public network 101.

[0010] Distributed data processing system 100 optionally has its own private communications networks 102. Communications on network 102 can be done through a number of means: standard networks just as in 101, shared memory, shared disks, or anything else. In the depicted example, a number of servers 110, 111, 112, and 113 are connected both through the public network 101 as well as private networks 102. Those servers make use of the private network 102 to reduce the communication overhead resulting from heartbeating each other and running membership and n-phase commit protocols.

[0011] In the depicted example, all servers are connected to a shared disk storage device 124, preferably a Redundant Array of Independent Disks (RAID) device for better reliability, which is used to store user application data. Data are made highly available in that when a server fails, the shared disk partition and logical disk volume can be failed over to another node so that data will continue to be available. The shared disk interconnection can be a Small Computer System Interface (SCSI) bus, Fibre Channel, or IBM Serial Storage Architecture (SSA). Each server machine can also have local data storage device 120, 121, 122, and 123.

[0012] (It will be appreciated that the configuration in Figure 1 is intended purely as an example, and not as an architectural limitation for range of applicability of the present invention).

[0013] Referring now to Figure 2a, cluster computer system 200 using Microsoft Cluster Services (MSCS) is designed to provide high availability for NT Server-based applications. The initial MSCS supports failover capability in a two-node 202, 204, shared disk 208 cluster.

[0014] Each MSCS cluster consists of one or two nodes. Each node runs its own copy of Microsoft Cluster Services. Each node also has one or more Resource Monitors that interact with the Microsoft Cluster Services. These monitors keep the Microsoft Cluster Services "informed" as to the status of individual resources. If necessary, the resource Monitor can manipulate individual resources through the use of Resource DLLs. When a resource fails, Microsoft Cluster Services will either restart it on the local node or move the resource group to the other node, depending on the resource restart policy and the resource group failover policy and cluster status.

[0015] The two nodes in a MSCS cluster heartbeat 206 each other. When one node fails, i.e., fails to send heartbeat signal to the other node, all its resource groups will be restarted on the remaining node. When a cluster node is booted, the cluster services are automatically started under the control of the event processor. In addition to its normal role of dispatching events to

other components, the event processor performs initialization and then tells the node manager, also called the membership manager, to join or create the cluster.

[0016] The node manager's normal job is to create a consistent view of the state of cluster membership, using heartbeat exchange with the other node managers. It knows who they are from information kept in its copy of the cluster configuration database, which is actually part of the Windows NT registry (but updated differently). The node manager initially attempts to contact the other node, and if it succeeds, it tries to join the cluster, providing authentication (password, cluster name, its own identification, and so on). If there's an existing cluster and for some reason our new node's attempt to join is rebuffed, then the node and the cluster services located on that node will shutdown.

[0017] However, if nobody responds to a node's requests to join up, the node manager tries to start up a new cluster. To do that, it uses a special resource, specified like all resources in a configuration database, called the quorum resource. There is exactly one quorum resource in every cluster. It's actually a disk; if it is, it's very preferable to have it mirrored or otherwise fault tolerant, as well as multi-ported with redundant adapter attachments, since otherwise it will be a single point of failure for the cluster. The device used as a quorum resource can be anything with three properties: it can store data durably (across failure); the other cluster node can get at it; and it can be seized by one node to the exclusion of all others. SCSI and other disk protocols like SSA and FC-AL allow for exactly this operation.

[0018] The quorum resource is effectively a global control lock for the cluster. The node that successfully seizes the quorum resources uniquely defines the cluster. The other node must join with that one to become part of the cluster. This prohibits the problem of a partitioned cluster. It is possible for internal cluster communication to fail in a way that brakes the cluster into two parts that cannot communicate with each other. The node that controls the quorum resource is the cluster, and there is no other cluster.

[0019] Once a node joins or forms a cluster, the next thing it does is update its configuration database to reflect any changes that were made while it was away. The configuration database manager can do this because, of course, changes to that database must follow transactional semantics consistently across all the nodes and, in this case, that involves keeping a log of all changes stored on the quorum device. After processing the quorum resource's log, the new node will begin to acquire resources. These can be disks, IP names, network names, applications, or anything else that can be either off-line or on-line. They are all listed in the configuration database, along with the nodes they would prefer to run on, the nodes they can run on (some may not connect to the right disks or networks), their relationship to each other, and everything else about them. Resources are typically formed into and managed as

resource groups. For example, an IP address, a file share (sharable unit of a file system), and a logical volume might be the key elements of a resource group that provides a network file system to clients; Dependencies are tracked, and no resource can be part of more than one resource group, so sharing of resources by two applications is prohibited unless those two applications are in the same resource group.

[0020] The new node's failover manager is called upon to figure out what resources should move (failover) to the new node. It does this by negotiating with the other node's failover managers, using information like the resources' preferred nodes. When they have come to a collective decision, any resource groups that should move to this one from the other node are taken off-line on that node; when that is finished, the Resource Manager begins bringing them on-line on the new node.

[0021] Every major vendor of database software has a version of their database that operates across multiple NT Servers. IBM DB2 Extended Enterprise Edition runs on 32 nodes. IBM PC Company has shipped a 6-node PC Server system that runs Oracle Parallel Servers. There is no adequate system clustering software for those larger clusters.

[0022] In a 6-node Oracle Parallel Server system, those six nodes share the common disk storage. Oracle uses its own clustering features to manage resources and to perform load balancing and failure recovery. Customers that run their own application software on those clusters need system clustering features to make their applications highly available.

[0023] Referring to Figure 2B, DB2 typically uses a share nothing architecture 210 where each node 212 has its own data storage 214. Databases are partitioned and database requests are distributed to all nodes for parallel processing. To be highly available, DB2 uses failover functionality from system clustering. Since MSCS supports only two nodes, DB2 must either allocate a standby node 216 for each node 212 as shown. Alternatively, DB2 can allow mutual failover between each pair of MSCS nodes as shown in Figure 2c. In other words, two nodes 212, 212a are mutually coupled to two data storages 214, 214a. The former doubles the cost of a system and the latter suffers performance degradation when a node fails. Because database access is distributed to all nodes and are processed in parallel, the node that runs both its DB2 instance and the failed over instance becomes the performance bottleneck. In other words, if node 212a fails, then node 212 assumes its responsibilities and accesses data on both data storages, but runs its tasks in parallel.

[0024] Accordingly, the present invention provides a method of managing a clustered computer system having at least one node, said method comprising the steps of:

establishing a multi-cluster comprising said at least one node and at least one shared resource;

managing said at least one node with a cluster services program;

returning said system to an initial state after a failover event.

[0025] A preferred embodiment of the present invention provides a method for managing clustered computer systems and extends clustering to very large clusters by providing a mechanism to manage a number of cluster computer systems, also referred to as "clusters". In particular, the preferred embodiment detects an initiation of a restart of a cluster computer system within a number of cluster computer systems. The initiation of the restart of the cluster computer system will cause the cluster computer system to restart in a selective state; In addition, this cluster computer system includes one or more resources. In response to a determination that one or more of the resources within the cluster computer system that is being restarted is presently operating in another cluster computer system within the cluster computer systems, the restart of these resources will be prevented.

[0026] The improved method for managing a cluster computer system of the preferred embodiment of the invention supports a failover from one node to another node chosen from a group of many nodes.

[0027] The invention further provides a distributed data processing system having at least one node, and including:

means for establishing a multi-cluster comprising said at least one node and at least one shared resource;

means for managing said at least one node with a cluster services program;

means for returning said system to an initial state after a failover event.

[0028] The invention further provides a computer program including instructions executable by a distributed data processing system for performing a method of managing a clustered computer system having at least one node, said method comprising the steps of:

establishing a multi-cluster comprising said at least one node and at least one shared resource;

managing said at least one node with a cluster services program;

returning said system to an initial state after a failover event.

[0029] Thus although the description hereafter will focus on a fully functioning data processing system, the computer program of the invention is typically capable of being distributed in the form of a computer readable medium of instructions in a variety of forms, including recordable-type media such a floppy disc, a hard disk drive, a RAM, and CD-ROMs, and transmission-type

media such as digital and analog communications links.

[0030] It will be appreciated that the system and computer program of the invention enjoy the same preferred features as the method of the invention.

[0031] Viewed from another aspect, the invention provides a method of managing a clustered computer system having at least one node, said method comprising the steps of:

- (a) establishing a multi-cluster comprising said at least one node and at least one shared resource;
- (b) managing said at least one node with a cluster services program; wherein said cluster services program manages using a resource API within the at least one node; including managing a heartbeat signal sent between said at least one node and any other node within the multi-cluster;
- (c) failing over between a first node and any other node within the multi-cluster;
- (d) updating a cluster wide data file; and
- (e) returning said system to an initial state after a failover event.

[0032] Viewed from another aspect, the invention further provides a method in a distributed data processing system for managing a plurality of cluster computer systems, the method comprising:

detecting an initiation of a restart of a cluster computer system within the plurality of cluster computer systems, wherein the cluster computer system will restart in a selected state and includes a resource; and responsive to a determination that the resource is presently operating in another cluster computer system within the plurality of cluster computer systems, preventing a restart of the resource in the cluster computer system.

[0033] Viewed from another aspect, the invention further provides a distributed data processing system, having a plurality of cluster computer systems, comprising:

detection means for detecting an initiation of a restart of a cluster computer system within the plurality of cluster computer systems, wherein the cluster computer system will restart in a selected state and includes a resource; and preventing means, responsive to a determination that the resource is presently operating in another cluster computer system within the plurality of cluster computer systems, for preventing a restart of the resource in the cluster computer system.

[0034] A preferred embodiment of the invention will now be described in detail by way of example only with reference to the following drawings:

Figure 1 is a pictorial representation of an exemplary distributed data processing system in which the present invention may be implemented;

Figures 2a, 2b, and 2c provide illustrations of the Microsoft Wolfpack product;

Figures 3, 3a, 3b, 3c, and 3d illustrate a preferred embodiment of the present invention having an implementation across multiple clusters such as MSCS clusters;

Figures 4, 4a, and 4b are flowcharts of methods used in a preferred embodiment of the present invention to control multiple clusters; and

Figures 5 and 6 are SQL tables containing example configuration, status, and event processing rules used in a preferred embodiment of the present invention.

[0035] The approach described herein extends the Microsoft Cluster Manager functionality to manage a larger cluster but otherwise preserves its ease-of-use characteristics. When discussed in this application, a "multi-cluster" refers to a cluster of two or more cluster computer systems.

[0036] The present cluster system supports resource group failover among any two nodes in a larger cluster of two or more nodes. The present system also preserves the application state information across the entire cluster in the case of failure events. Also, the present system does not require change implementation of currently available cluster computer system products. For example, with respect to MSCS, the mechanism does not require Microsoft and application vendors to make any modification to their present clustering code in order to run in this system's environment. Instead, the present system provides an implementation of the MSCS Cluster API DLL that is binary compatible with the MSCS Cluster API DLL.

[0037] A multi-cluster normally contains more than one pair of clusters. A cluster manager in accordance with the preferred embodiment of the invention can configure a cluster with multiple MSCS clusters within. Resources in a multi-cluster are managed by each individual cluster under the supervision of Cluster Services. No need exists to modify the Microsoft Resource API and the Microsoft Cluster Administrator extension API. The Cluster Manager can use any Cluster Administrator Extension DLL that is developed for MSCS as it is without modification.

[0038] Applications, whether they are enhanced for MSCS or not, can readily take advantage of system clustering features described herein. Instead of mutual failover between one pair of nodes, an application failover is allowed between any two nodes in a large cluster. This allows a cluster to grow in size by adding an MSCS cluster either with a pair of nodes or a single node. The fact that a three node cluster can be supported is very attractive to many customers who want to further improve availability of their mission critical appli-

cations over a two node cluster.

[0039] Applications such as DB2 Extended Enterprise Edition that use MSCS can readily take advantage of multi-cluster system clustering features. DB2/EEE exploits MSCS features by dividing nodes into pairs and allows mutual failover between each pair of nodes as discussed above in reference to Figure 2c. The approach described herein can either improve DB2 availability by supporting N-way fail-over or improve DB2 performance characteristics by supporting N+1 model with one standby node. In the most common event of a single node failure, DB2/EEE instance on the failed node will be restarted on the standby node and maintain the same performance in the N+1 mode. System management policy and recovery services are expressed in a high-level language that can be modified easily to tailor to special requirements from application vendors. For example, this allows DB2/EEE to be integrated with a multi-cluster better than with a MSCS cluster.

[0040] It must be understood that the approach described herein can be used over any cluster service program. While the depicted example illustrates MSCS clusters within a multi-cluster, the processes, mechanisms, and instructions may be applied to managing clusters of all types. Thus applicability is not limited in any way to use over that particular product, and, for example, potentially extends also to heterogeneous multi-clusters.

[0041] With reference now to Figure 3, a pictorial representation of a distributed data processing system is depicted. The software 300 shown in Figures 3, 3b, and 3c can be implemented on the hardware shown in Figure 3a. The process for the multi-cluster software illustrated herein can scale to larger sizes easily. For example, Figure 3a shows an eight-node configuration, wherein each node 350 is coupled to a storage element 340 by disk controllers 360. cluster services 304 (which represent the focus of the present invention) in Figure 3 allow fail-over to be between any two nodes in this eight-node cluster. The cluster services, such as cluster services 304, are employed to control a cluster, such as a MSCS cluster, and can be used in both the Oracle cluster or a DB2 cluster discussed above. In the case when any of the seven nodes fails, the DB2 instance will be restarted on the eight node and the performance of the system will remain unchanged. This is called an N+1 failover model. Other configurations are also supported. For example each node may run an active DB2 instance and be backup for the other seven nodes to maximize reliability.

[0042] MSCS is used to perform resource management for a single node in the depicted example. Microsoft does not share its resource management APIs in Windows NT with outside vendors and there is no easy way for other vendors to perform resource management. Some vendors have implemented their own device drivers and TCP/IP protocol stack. That results in

incompatibility with the MSCS Cluster API and Resource API. The present approach uses MSCS to manage resources on a single node, and thus does not need to know the internal NT APIs. Again, while reference is made herein to the Microsoft cluster product, there is no limitation of the approach described herein to use over that product, rather it can be used over any suitable cluster services program.

[0043] Referring to Figure 3, cluster services 304 controls MSCS 306 to bring a resource and a resource group on-line or off-line on a node 350. Cluster services 304 is shown controlling the MSCS 306 and 306a, which are located on different nodes 350 and 350a. Cluster Services 304 causes MSCS 306 to bring resource group containing application 370 off-line and then cause MSCS 306a to bring that resource group on-line. Cluster services 304 is responsible for managing cluster node membership, heartbeat, inter-node communications, and for maintaining the consistency of cluster configuration database for all eight nodes. Cluster services also is responsible for event notification and processing. Cluster manager 302 provides a graphical user interface (GUI).

[0044] Cluster services 304 is substantially binary compatible with MSCS in this example. No modification is required to run any application in a multi-cluster if that application can run in an MSCS cluster. Cluster services supports all MSCS Cluster API, Resource API, and Administrator Extension API.

[0045] Referring to Figures 3b and 3c, in a multi-cluster, each node runs a copy of Cluster Services. When a node 350 is booted, cluster services 304 is started automatically. The MSCS cluster services 306 is then started by cluster services 304. In this document, we will refer to those MSCS clusters within a multi-cluster as MSCS subclusters. The configuration information in a multi-cluster configuration database is a super set of the information in each MSCS subcluster. All resources and resource groups are defined in a multi-cluster configuration database and in appropriate MSCS subclusters. When an MSCS subcluster services is started, all resources and resource groups except the default Cluster Group are left in an off-line state. Cluster services 304 on a new node determines collectively through CSQL\_Services group 315 with cluster services instances on all other nodes which resource groups should be started on that node. It then invokes the MSCS cluster services API to bring those resource groups to an on-line state.

[0046] Each MSCS subcluster consists of either a pair of nodes or a single node. In the case of single-node MSCS subcluster, the MSCS quorum resource can be configured as a local quorum resource, which means that the quorum resource will be a local disk of that node. This is a preferred configuration since it will save a shared disk per MSCS subcluster.

[0047] Some cluster servers, such as, for example, MSCS, have a unique feature in that they remember the

state of resources and resource group the last time when the cluster was terminated. When a node is restarted, MSCS cluster services will bring those resources and resource groups to the previous state. The decisions regarding bringing resources and resource groups to their on-line and off-line state are made by the multi-cluster services. If an MSCS subcluster (or the node that runs that MSCS subcluster) fails, the cluster services will restart those resources and resource groups that were running on that node on some other MSCS subcluster. When the failed node and the corresponding MSCS subcluster is restarted and re-joins the multi-cluster, there will be resource conflicts if the new node and new MSCS subcluster try to bring those resources and resource groups to an on-line state. To resolve this problem, cluster services adds a "hidden" resource into every resource group and makes this hidden resource a dependent resource for all other resources in that resource group. The hidden resource will check the state of its resource group in the multi-cluster configuration database and will fail to start if the resource group is already running on another MSCS subcluster.

[0048] The cluster services extends the high availability system clustering features of presently available cluster services to more than two nodes and preserves binary compatibility with presently available cluster services.

[0049] Referring now to Figures 3b and 3c, the present system clustering software 300 consists of two major parts: cluster manager 302 and the cluster services 304. The cluster manager 302 is designed to manage all resources in a group of clusters 306 and to present a single cluster image to its users. The cluster manager 302 provides an easy-to-use user interface that information technology (IT) administrators are accustomed to. The cluster manager 302 allows administrators to manage a large scale and complex collection of highly available resources in a cluster efficiently and effectively.

[0050] The cluster services 304 is a middleware layer that runs on each computer 350 in the cluster. In the depicted example, it comprises a set of executables and libraries that run on the resident Microsoft Windows NT server or other suitable server. The cluster services 304 contains a collection of inter-acting subsystems. Those subsystems are Topology Services 308, Group Services 310, cluster Coordinator (not shown), CSQI Services 314, Event Adapters 310, Recovery Services 316, and the cluster API 318.

[0051] The Cluster Coordinator provides facilities for start up, stop, and restart of cluster services 304. There is a Cluster Coordinator on each computer in the cluster, but they do not communicate with each other; each one's scope is restricted to the computer on which it runs. The Cluster Coordinator is the component that needs to be started up first. It then brings up the other services in the following order: CSQI Services 314 in

stand-alone mode; Topology Services 308; Group Services 308; CSQI Services 314 in Cluster-mode; Recovery Services 316; Microsoft Cluster Services (MSCS) Event Adapter; MSCS; and Group Services Event Adapter (GSEA). Further, it monitors each of the other services, and terminates all other services and user applications and restarts the multi-cluster cluster services in case of failures.

[0052] Topology Services 308 sends special messages called heartbeats that are used to determine which nodes are active and running properly. Each node checks the heartbeat of its neighbor. Through knowledge of configuration of the cluster and alternate paths, Topology Services 308 can determine if the loss of a heartbeat represents an adapter failure or a node failure. The MSCS's inter-node heartbeat is ignored in favor of the topology services heartbeat which is multi-cluster wide. Topology Services maintains information about which nodes are reachable from which other nodes, and this information is used to build a reliable messaging facility.

[0053] Group Services 310 allows the formation of process groups containing processes on the same or different machines in the cluster. A process can join a group as a provider or a subscriber. Providers participate in protocol actions, discussed in detail below, on the group while subscribers get notified on changes to the group's state or membership (list of providers). Group Services 310 supports notification on joins and leaves of processes to a process group. It also supports a host group that one can subscribe to in order to obtain the status of all the nodes in the cluster. This status is a consistent view of the node status information maintained by Topology Services.

[0054] All MSCS subclusters in a multi-cluster are preferably configured as single-node clusters. Group Services is used for monitoring node up and node down events. Group Services also provides the following facilities for cluster-aware applications to handle failure and reintegration scenarios. These facilities are built on top of the reliable messaging facility: Atomic broadcast and n-phase commit protocols for process join, process leave - voluntary and involuntary, process expel, group state change, and provider broadcast messages

[0055] Group Services 310 handles partitioning of the cluster in the following manner. When it recognizes that a cluster that was partitioned has come together, it will generate a dissolve notification to all groups that were part of the partition that has the lesser number of cluster machines. If both partitions have equal number of cluster machines, one of them is chosen to be dissolved.

[0056] CSQI Services 314 provides support for a database that can contain configuration and status information. It can function in both stand-alone and cluster modes. Each database is a persistent, distributed resource which, through the use of Group Services 310, is guaranteed to be coherent and highly available. Each

database is replicated across all nodes and check-pointed to disk so that changes are obtained across reboots of the multi-cluster cluster services. CSQL Services 314 ensures that each node has an identical copy of data. CSQL Services also supports a transient type of data that does not persist across reboot but is also consistent on all nodes. Transient data will be initialized to their startup values after a restart of cluster services 304. CSQL Services 314 supports notification of changes made to the database. Each database can be marked by a three tuple: a timestamp indicating when a database was last modified, the ID of the node that proposed the modification, and a CRC checksum. The timestamp is a logical time that is a monotonically increasing number across the entire cluster. CSQL Services 314 runs a Database Conflict Resolution Protocol to determine the most up-to-date replica upon a cluster restart. A node replaces its replica by the cluster's version after making a backup of the existing version of each replace databasa when it rejoins a cluster. Modification to a cluster configuration database is permitted only after CSQL transits from stand-alone mode to cluster mode. The conditions for entering cluster mode will be discussed thoroughly below. CSQL Services supports both local and remote client connections.

[0057] Event Adapter 312 monitors conditions of subsystems, and generates events when failure situations occur. Events are inserted into a distributed event queue, which is implemented as an event table in the cluster-scope CSQL configuration database. There are four event adapters in a cluster: MSCS Event Adapter that monitors the MSCS subsystem, Group Service Event Adapter that monitors node and network interface failures, Cluster API Event Adapter that converts user requests into multi-cluster events, and Partition Prevention Event Adapter that monitors network partition.

[0058] Group Services Event Adapter (GSEA) 310 is a distributed subsystem. Each GSEA joins a GSEA Group Services group 311 as a provider. GSEA receives LEAVE and FAILURE LEAVE notification from Group Services and converts them into multi-cluster events. GSEA as a group inserts exactly one event into the event queue when a GSEA leaves the group either voluntarily or due to failure.

[0059] Microsoft Cluster Services Event Adapter (MSCSEA) 320 converts a MSCS notification into events recognizable by the present cluster manager. There is one instance of MSCSEA running on each node. Each MSCSEA is used to monitor MSCS resource groups and MSCS resources that are running on the local node only. When MSCS subclusters in a multi-cluster are configured as single-node clusters and therefore the MSCS heartbeat mechanism is effectively disabled, network interface failure and node failure will be detected by the Topology and Group Services subsystem 308.

[0060] Recovery Services 316 is a rule-based object-oriented, and transactional event processing

subsystem. Event processing is triggered when a new event is inserted into the cluster-wide event table in a cluster-scope CSQL database. Recovery Services extends the CSQL functionality and added active and object-oriented SQL statement processing capability into the CSQL subsystem. Methods are expressed in the active SQL language. Specifically, the following SQL-like active SQL statements are introduced: CREATE TRIGGER, EVALUATE, EXECUTE, CONTINUE, CREATE MACRO, and LOAD DLL. The CREATE TRIGGER statement registers a trigger on the specified table with CSQL. When a new row (event) is inserted into the specified table, CSQL will invoke the corresponding event processing rules. Rules are expressed in SQL and the above mentioned active SQL statements. EVALUATE statement is very similar to SELECT. Instead of selecting a set of data, an EVALUATE statement selects a set of rules and then evaluates those rules. SQL and active SQL statements that are selected and processed by the same EVALUATE statement are part of the same transaction. The EXECUTE statement changes the physical system state by invoking either a user defined function, an external program, a command file, or a shell script file. The CONTINUE statement synchronizes event processing among distributed CSQL Servers. In particular, CONTINUE statement synchronizes the CSQL database until the point of the CONTINUE statement. There can be multiple CONTINUE statements each time when event processing is triggered. The Create MACRO statement defines the specified macro, which can be invoked in any SQL statement. A macro returns a data value that can be used in a SQL statement. LOAD DLL dynamically loads the specified dynamically linked library (DLL) into CSQL. During the DLL initialization code, it registers those user defined functions in the DLL into CSQL. User defined functions can be invoked either in an EXECUTE statement or embedded in any other SQL statement. User defined function extends SQL language either by providing commonly used functionality or by initiating actions on physical entities external to CSQL Server. As an example, user defined functions are used to control MSCS resource management facilities.

[0061] Although one preferred embodiment of the cluster services for a multi-cluster is shown, other mechanisms may be used to provide cluster services. For example, CSQL programming interface takes SQL statements. Other types of programming interfaces or data storage or data registration mechanisms may be employed. In such an implementation, the mechanism would provide consistency of data across the clusters within the multi cluster, provide consistency of data for the various nodes during a reboot, and provide synchronization of data for a new node entering a cluster. In addition, although the recovery services described in the preferred embodiment are an extension of CSQL, other embodiments may be constructed that do not require such an extension.



[0062] Multi-cluster API 318 provides access to a multi-cluster as a whole, not a particular MSCS cluster. It contains functions that can handle a larger cluster but otherwise is functionally identical to those functions of the Microsoft Cluster API. It is intended to be used by Cluster Manager 302 as well as other cluster-aware applications. There is a one-to-one correspondence between functions in the Multi-Cluster API and those in the Microsoft Cluster API. The similarity between the two Cluster APIs can help application vendors take advantage of multi-clustering features now and to migrate to greater-than-two-node Microsoft clusters in the future. The Multi-Cluster API DLL is binary compatible with the MSCS Cluster API DLL, clusapi.dll. The query type of Cluster API functions are handled directly by the Multi-Cluster API DLL. Those Cluster API functions that cause state changes are converted into events which are handled by Recovery Services. Multi-Cluster API DLL uses CSQL Notification to wait for the result of event processing. Multi-cluster API DLL communicates with CSQL Services via a well-known virtual IP address. In sum, the cluster services 304 guarantees that the state information put into the NT cluster registry by an application program will be available when that application falls over to another node in a cluster. The cluster services 304 provides utilities that examine the system configuration and make sure that a system is properly configured for installation and running system clustering features. Clusters are configured accordingly when it is first started. Accompanying cluster services 304, the Cluster Manager 302 will configure, manage, and monitor clusters and their contained MSCS clusters. Other utilities may be developed to help simplify the installation process of multiple MSCS subclusters and the multi-cluster cluster services.

[0063] The cluster services subsystems are started by the Cluster Coordinator subsystem. The Cluster Coordinator is implemented as an NT service and is started automatically during startup. The cluster coordinator then starts all other Cluster Services subsystems in the following order: CSQL Services in stand-alone mode, Topology Services, Group Services, CSQL Services in cluster mode, Recovery Services, MSCS Event Adapter, MSCS, and Group Services Event Adapter.

[0064] CSQL Services is initially started in stand-alone mode. Topology Services and Group Services retrieve their configuration information from CSQL databases. After Group Services comes up, CSQL Services forms the CSQL\_Services group 315 and runs a Database Conflict Resolution Protocol (DCRP) to synchronize the contents of the cluster configuration database. The first CSQL server forms the group, sets the CSQL\_Services group in a BIDDING state, and starts a timer to wait for other CSQL servers to join the group. A CSQL server that joins the group which is in the BIDDING state also starts a timer to wait for others to join. The timer value is defined in the cluster configuration database and may be different from node to node.

Inconsistent timer values can be caused by different versions of cluster configuration databases that are being used by different nodes initially. When the first timer expires, the CSQL server broadcasts the timestamp of its cluster configuration database to the group using a Group Services n-phase protocol. Other CSQL servers broadcast their timestamps if their timestamp is more recent than the received one. When multiple CSQL servers send out their timestamp, one will be selected arbitrarily by Group Services and broadcast to the group in the next phase. A CSQL server sends out its timestamp only if its timestamp is better than the received timestamp. A CSQL server should send out its timestamp even if it is older than the received one only in the first phase in order to signal other CSQL servers that it has a different version. Eventually the protocol will conclude. Either all CSQL servers have an identical timestamp or they all agree on the most up-to-date version. If not all timestamps are identical, the CSQL server that sends out its timestamp last should broadcast its database to all others. CSQL servers should make a backup copy for any database that is to be replaced by the latest version. After CSQL servers synchronize the cluster configuration database, they will set the state of the CSQL\_Services group to its RUNNING state. Those CSQL Servers whose replica was replaced by a new version will initiate a restart of Cluster Services. A CSQL server that joins a RUNNING CSQL\_Services group must save its replica and replace it by the cluster version regardless of its timestamp value. If the new version has a different timestamp from the existing one which is presently being used by other subsystems, the CSQL Server will initiate a restart of Cluster Services.

[0065] The CSQL timestamp is a three tuple: a monotonically increasing number across the entire cluster, the node ID of the node that modified the database the last time, and a CRC check sum.

[0066] Once CSQL Services is in a RUNNING state, the cluster configuration database including the event queue is consistent on all nodes. A CSQL server is said to be in cluster mode after it successfully joins a RUNNING CSQL\_Services group. Recovery Services, MSCS, MSCS Event Adapter (MSCSEA), and Group Services Event Adapter (GSEA) will then be started. The GSEA joins a GSEA Group Services group and adds a BRING\_COMPUTER\_UP event for this node into the cluster-wide event queue in processing the Group Services JOIN protocol. Multi-cluster resource groups are initially in an off line state. During the processing of a BRING\_COMPUTER\_UP event, Recovery Services determines whether any resource group should be brought into an online state.

[0067] The DCRP algorithm is summarized below: (1) A CSQL server broadcasts an open database request including the name of the database and a timestamp to the CSQL\_Services group, (2) Each CSQL server that has a different timestamp must vote

CONTINUE and broadcast its timestamp in the first phase to force a database replication, (3) The CSQL server that receives its own broadcast must vote APPROVE in the first phase, (4) A CSQL server that has identical timestamp as the received one must vote APPROVE, (5) For each subsequent phase, a CSQL server that has a later timestamp than the received one must broadcast its timestamp and vote CONTINUE, (6) A CSQL server that receives its own timestamp must vote CONTINUE, (7) A CSQL server that has the same or any earlier timestamp must vote APPROVE, (8) If no message was sent in a phase, the server that broadcast its timestamp last must replicate its version of the database to other servers. A server always makes a backup copy of its replica before replacing it.

[0068] Still referring to Figures 3b and 3c, the start-up sequence for the multi-cluster system is illustrated. First, the Cluster Coordinator is started as NT Services during NT startup. The Cluster Coordinator starts and monitors other multi-cluster subsystems. Next, CSQL Services 314 is started in stand-alone mode. Then, Topology Services 308 is started. Group Services 310 is then started. Next, CSQL Services forms or joins the CSQL\_Services group 315. CSQL Services runs the Database Conflict Resolution Protocol and enters cluster mode. Then all cluster scope databases are up-to-date. In particular, the event queue is up to date. Recovery Services 316 is started and Recovery Services daemon starts both the MSCS Event Adapter 312 and the group Services Event Adapter 310, in this order. Group Services Event Adapter (GSEA) 310 is started. GSEA forms or joins the GSEA group and it will monitor node failure events. Recovery Services daemon then inserts A BRING\_COMPUTER\_UP event for the local node. Recovery Services processes the BRING\_COMPUTER\_UP event for this node. MSCS subsystem 306 is started and then monitored by the MSCS Event Adapter 312. Resource groups are started or moved to this new node depending on resource allocation policy and system status.

[0069] Another important feature of the preferred embodiment of the present invention involves a cluster quorum condition. No resource group can be brought into its online state unless one of the following quorum conditions has been met. Cluster Services adopts the same majority quorum scheme that is used in HACMP. Cluster Services uses connectivity information provided by Group Services to determine majority quorum condition. Additionally nodes also pass connectivity information through the shared disk path or use some other method to avoid the split brain problem. When the network is severed and a cluster is divided into several partitions, Cluster services must guarantee not to start a single resource group in multiple partitions at the same time which can cause corruption to application data on shared disks. The connectivity information passed on the disk path helps each partition to learn about the

sizes of other partitions and hence help prevent data corruption. A resource group should be brought into the online state on one if the following conditions are true: (1) the partition has a majority quorum, i.e., more than half of all nodes defined in the cluster configuration database have joined a cluster and are in that partition, or (2) the partition has exactly half of the nodes as defined in the cluster configuration database and no other partitions of the same size exist, or (3) the partition has exactly half of the nodes as defined in the cluster configuration database while another partition contains the other half of the nodes and the smallest node ID is in the former partition.

[0070] After starting all Cluster Services subsystems, the Cluster Coordinator will monitor the status of each subsystem. If any subsystem terminates abnormally, the Cluster Coordinator will shutdown the node and will restart itself, as well as other subsystems. Shutting down a node when any subsystem fails can guarantee that no user applications will continue running when the Cluster Services fails.

[0071] When a partition heals, Group Services will resolve groups in all but one partition. Group Services daemon in those partitions will be terminated. Consequently those nodes will be shut down by the Cluster Coordinator and restarted. The shutdown procedure for Recovery Services must make sure that all resource groups are offline.

[0072] Referring to Figure 3C, the component support in the preferred embodiment is illustrated. Cluster services 304 uses MSCS 306 to manage cluster resources within a node. A resource group is defined in cluster configuration database first and defined in a MSCS subcluster only if needed. Resource management policy is designed to mimic the MSCS resource management behavior. When a resource group is defined in a MSCS subcluster, the restart flag is always disabled so that a restart decision will be made by the event processing subsystem, not by MSCS. A resource group defined in an MSCS subcluster, irrespective of whether it is a single node cluster, will have at most one node in the preferred node list so that the MSCS auto failover mechanism is disabled. Cluster services will monitor the status of every resource group that is online. When a resource or resource group failure occurs, the MSCS event adapter 312 will insert the corresponding event into the event queue. CSQL services 314 will trigger event processing for the event. One and only one CSQL instance will initiate event processing. Each CSQL instance manages resources including the single-node MSCS subcluster on the local node only. Event processing is designed to be able to handle multiple failures.

[0073] Referring now to Figures 4, 5, and 6, another aspect of the present system involves Event Processing. With respect to Figure 5, table 500 illustrates two entries 502 and 504, which describe two ch\_routines: BRING\_COMPUTER\_UP and NODE\_UP.

In entry 502, the action in section 506 corresponds to step 404 in Figure 4. In entry 504, sections 508, 510, and 512 contain actions that correspond to steps 408, 410, and 414, respectively. Events defined in Cluster services include but are not limited to: BRING\_COMPUTER\_UP, BRING\_COMPUTER\_DOWN, BRING\_RESOURCE\_GROUP\_ONLINE, BRING\_RESOURCE\_GROUP\_OFFLINE, and MOVE\_RESOURCE\_GROUP. When a computer joins a cluster, a "BRING\_COMPUTER\_UP" event will be inserted into the event queue. To process a BRING\_COMPUTER\_UP event, the cluster services performs the following: (1) Check whether a quorum exists, and (2) If so, then check whether any resource group should be brought up on the new computer. Some resource groups may be online on some other computer. Those resource groups should be brought into an off line state first. Next, the cluster services should bring those resource groups that are in an off line state online on the new computer.

[0074] All the configuration information, status information, resource management policy, and rules are stored in a cluster scope database, escluster.cfg. Suppose that computer "hilltop" joins a cluster. A BRING\_COMPUTER\_DOWN event for hilltop is inserted into the event queue, which triggers CSQL to perform event processing wherein a runtime environment is created which encapsulates the information relevant to the event and CSQL processes the following statement:

```
EVALUATE action from ch_routines where
ch_routine = "BRING_COMPUTER_UP"
```

[0075] The above statement specifies that statements in the BRING\_COMPUTER\_UP row of the ch\_routines table in the escluster.cfg database should be processed. The actions taken in a ch\_routine called BRING\_UP\_COMPUTER are depicted in table 500 in entry 502.

[0076] The ch\_resource\_groups table 600 is defined in Figure 6, which shows one row of the table, with each entry representing one column. \$failback\_node() is a macro which returns a node where the specified resource group should be running based on the specified failback policy and given the fact that a new node rejoins a cluster. \$resource\_group\_online() and \$resource\_group\_offline() are user defined functions that use MSCS Cluster API function calls to bring the specified resource group off line and online on the specified computer node. As a result of processing "EVALUATE action from ch\_routines where ch\_routine = "BRING\_COMPUTER\_UP"", the following statements are selected and then processed:

```
evaluate markup_action from computers where
```

```
computer + $get_event_node();
evaluate action from ch_routines where $has_
quorum90 and ch_routine = NODE_UP
```

[0077] The actions taken for the ch-routine called NODE\_UP are illustrated in entry 504 of table 500 in Figure 5. As a result of processing of the second EVALUATE statement, the following three statements are retrieved and then processed:

```
evaluate failback_action from ch_resource_groups
where current_node <> next_node;
evaluate release_action from ch_resource_groups
where current_node <> next_node;
evaluate acquire_action from ch_resource_groups
where current_node = "" and next_node =
$get_event_node();
```

[0078] Those three EVALUATE statements will each search for all ch\_resource\_group rows (object) in the ch\_resource groups table that meet the search condition. When a ch\_resource\_group row (object) is found, the specified action will be applied to that object. The failback\_action contains a single statement, which is:

```
update ch_resource_groups set next_node =
$failback node() where ch_resource group = this
ch_resource_group
```

[0079] In the above update statement, a macro failback\_node() is processed which returns a node that is the most preferred node for running the specified resource group given that a new node has just joined the cluster. The update statement stores the returned node name into the next\_node column. A macro name is prefixed by \$\_ to simplify parsing.

[0080] The current\_node column of a ch\_resource\_group object indicates the current node on which the ch\_resource\_group is running. The release\_action is processed for this ch\_resource\_group if the current\_node is different from the next node. If that is the case, the following statement is processed:

```
execute $_resource_group_offline()
```

[0081] Resource\_group\_offline() is a user defined function which in term calls the MSCS OfflineResource-Group() function to bring the implied resource group to its off line state. A user defined function is prefixed by \$\_ to simplify parsing.

[0082] Finally, the acquire\_action is retrieved and processed on the new node for all those ch\_resource\_group objects that are not running anywhere and that should be running on the new node. The acquire\_action contains one statement:

```
execute $_resource_group_online()
```

[0083] Resource\_group\_online () is also a user defined function which calls the MSCS OnlineResourceGroup() function to bring the implied resource group to its online state.

[0084] Cluster Services also supports event simulation. When Recovery Services is invoked to simulate an event, it first clones the cluster configuration database. The event simulation will be performed on the private copy of the configuration database so that the original configuration database will not be affected. During a simulation, it is the EXECUTE statement which actually changes the state of physical resources.

[0085] Figure 4 illustrates the method implemented by the cluster services when a node wants to join a cluster. First, a node joins the cluster (step 402). A decision is made as to whether a quorum exists (step 404). If not, the method returns (step 406). If a quorum does exist, then for every resource group, the following loop is implemented (step 405). First a query is made whether any resource group should be failback to the new node (step 408). If so, then for each such resource group, the system gets the corresponding MSCS sub-cluster to do an off-line of the specified resource group (step 410). A continue (step 418) is performed to synchronize all the nodes. The MSCS sub-cluster on the new node will bring the specified resource group to the online state (step 414). A query is then made (step 412) to see if there are more resource groups. If not, the system is done (step 416); otherwise the method returns to step 405.

[0086] Figure 4a illustrates a flowchart of the method 430 to move a resource group from one node to another. Every node computes the next most preferred node to run the resource group based on node status, the resource group preferred node list, and the failover policy (step 434). Alternatively, the user can simply specify the next node. Next, the system queries if the current node is not equal to the next node (step 436). If not, the system is done (step 438). If so, then the system gets the MSCS sub-cluster on the current node to bring the specified resource group to off line (step 440). The process then continues (step 442). During this step, the system synchronizes its event processing. Afterwards, the system gets the MSCS cluster on the next node to bring the specified resource group to the online state (step 444). Finally, the system is done (step 446).

[0087] Figure 4b illustrates the general method 450 implemented by cluster services when node failure occurs. This method can also be applied to resource failure and resource group failure events. The group service event adapter collectively inserts exactly one node down event into the event queue (step 454). Node\_Down event processing is triggered (step 456). Next, for every resource group that was running on the failed node, the following steps are applied (step 458). First, recovery services compute the Next\_Node for failover (step 460). Then a decision is made if My\_Node == Next\_Node. If not, the system checks if there are

more resource groups (step 462). If so, then the system gets the MSCS sub-cluster to bring the specified resource group online (step 464). If no more resource groups are available, then the system is done (step 466). If more are available, then the system loops back to step 458.

[0088] While the preferred embodiment has been described as using MSCS sub-clusters, it is important to understand that many other different embodiment are possible. For example, an analogous system could be built on top of the IBM HACMP or the Sun Microsystems Ultra Enterprise Cluster HA Server to manage these cluster systems, or applied to heterogeneous clusters systems, such as for managing a multi-cluster system including a cluster managed using MSCS and a cluster using an Ultra Enterprise Cluster HA server. In addition, the approach described herein may be applied to managing multiple processor computers, such as SMP servers.

#### Claims

1. A method of managing a clustered computer system having at least one node, said method comprising the steps of:
  - establishing a multi-cluster comprising said at least one node and at least one shared resource;
  - managing said at least one node with a cluster services program;
  - returning said system to an initial state after a failover event.
2. The method of Claim 1, wherein said cluster services program manages using a resource API within the at least one node.
3. The method of Claim 1 or 2, wherein the multi-cluster comprises at least two clusters, wherein each of said clusters comprises at least one node.
4. The method of Claim 1 or 2, wherein the multi-cluster comprises at least three nodes.
5. The method of any preceding Claim, further comprising the step of failing over between a first node and any other node within the multi-cluster.
6. The method of Claim 5, further comprising the step of updating a cluster wide data file.
7. The method of any preceding Claim, wherein said managing step includes initiating a first cluster services program automatically when said at least one node is booted.
8. The method of Claim 7, wherein said managing

step further includes initiating a second cluster services program resident on the at least one node after initiating the first cluster services program.

9. The method of Claim 8, wherein said first and second cluster services programs are binary compatible. 5
10. The method of any preceding Claim, further comprising the step of managing a cluster node membership database. 10
11. The method of any preceding Claim, further comprising the step of sending a heartbeat signal between said at least one node and any other node within the multi-cluster. 15
12. The method of any preceding Claim, wherein said managing step includes managing inter-node communications between said at least one node and any other node within the multi-cluster. 20
13. The method of any preceding Claim, further comprising the step of presenting an image of a single cluster with a cluster manager. 25
14. The method of any preceding Claim, wherein said managing step includes configuring a multi-cluster quorum resource as a local quorum resource. 30
15. The method of any preceding Claim, wherein said returning step includes restarting a node and bringing said shared resource to the initial state.
16. The method of any preceding Claim, wherein said returning step includes storing said initial state for said shared resource. 35
17. The method of any preceding Claim, wherein said managing step includes, responsive to a conflict for control of said shared resource, failing to restart a failed node which would attempt to control said shared resource. 40
18. The method of Claim 17, further comprising the step of adding a hidden resource into a resource group on each node. 45
19. The method of Claim 18, further comprising the step of making said hidden resource dependent on any other resource in said resource group. 50
20. The method of any preceding Claim, wherein the distributed computer system includes a plurality of cluster computer systems, further comprising the steps of: 55

detecting an initiation of a restart of a cluster

computer system within the plurality of cluster computer systems, wherein the cluster computer system will restart in a selected state and includes a resource; and

responsive to a determination that the resource is presently operating in another cluster computer system within the plurality of cluster computer systems, preventing a restart of the resource in the cluster computer system.

21. The method of claim 20, wherein the resource is a shared file system.

22. A distributed data processing system having at least one node, and including:

means for establishing a multi-cluster comprising said at least one node and at least one shared resource;

means for managing said at least one node with a cluster services program;

means for returning said system to an initial state after a failover event.

23. A computer program including instructions executable by a distributed data processing system for performing a method as claimed in any of claims 1 to 21.




FIG. 1





FIG. 2a  
(PRIOR ART)



STANDBY'S

FIG. 2b  
(PRIOR ART)




FIG. 2c  
(PRIOR ART)

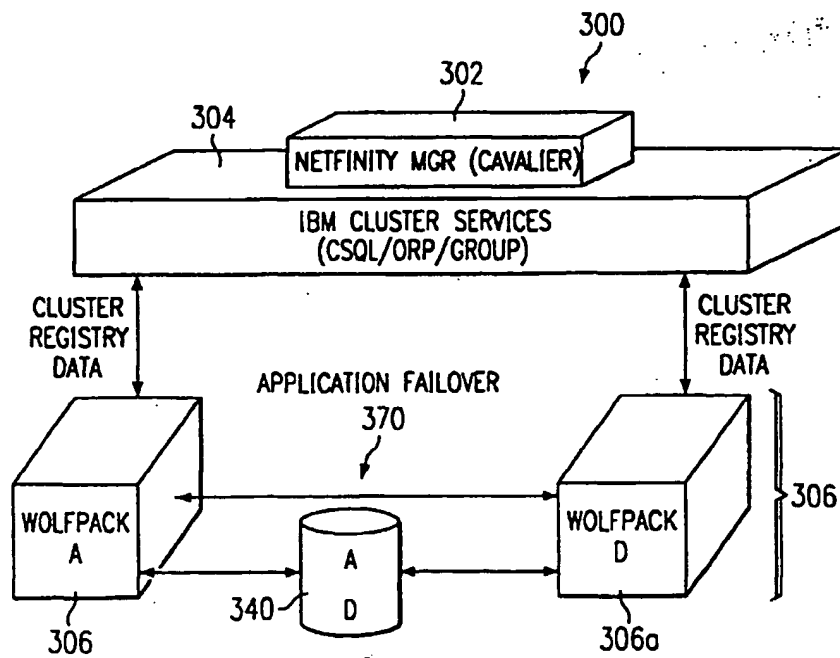



FIG. 3









FIG. 3b



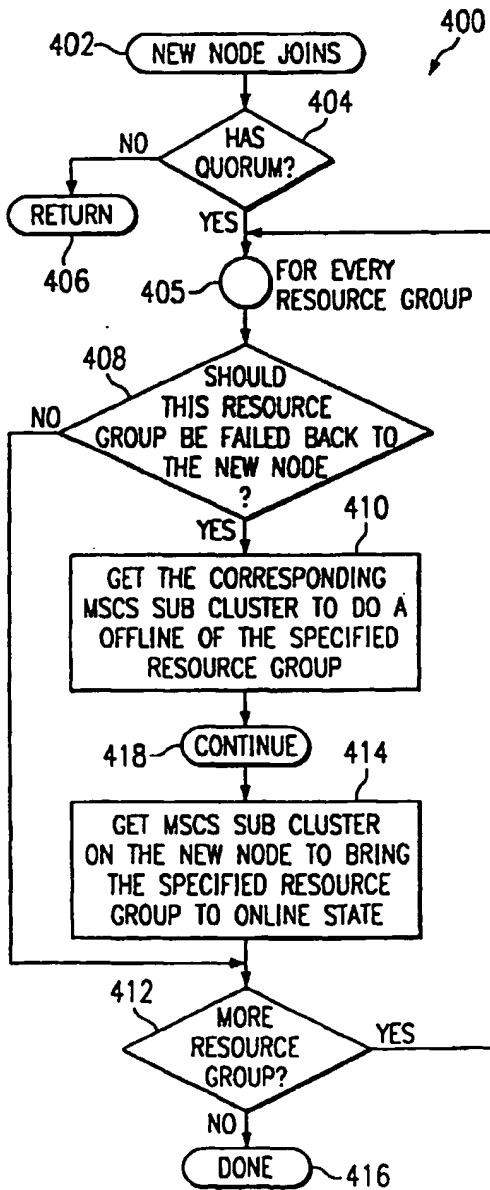


FIG. 4

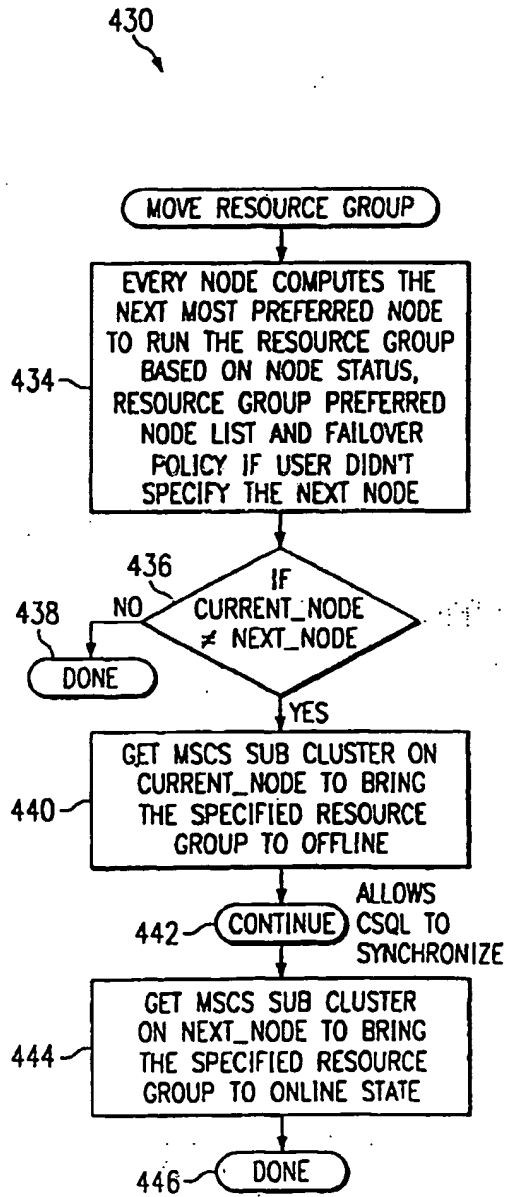


FIG. 4a

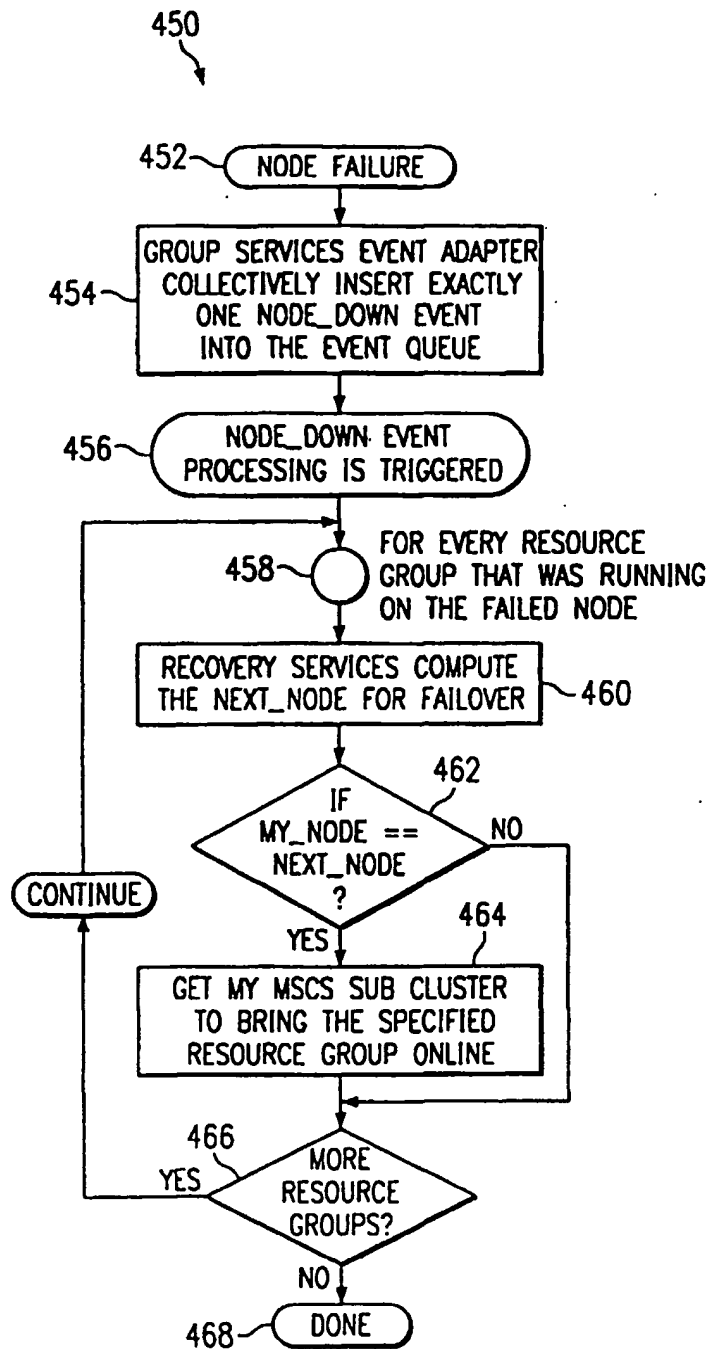


FIG. 4b

Ch routine	Action	500
BRING_COMPUTER_UP	Evaluate markup_action from computers where computer = \$_get_event_node() ; evaluate action from ch_routines where \$_has_quorum() and ch_routine = NODE_UP;	
NODE_UP	Evaluate failback_action from ch_resource_groups where current_node<>\$_get_event_node() ; evaluate release_action from ch_resource_groups where current_node <>next_node; evaluate acquire_action from ch_resource_groups where current_node = "" and next_node = \$_get_event_node() ;	<div data-bbox="1104 378 1161 462">501</div> <div data-bbox="1104 567 1161 609">508</div> <div data-bbox="1104 693 1161 735">510</div> <div data-bbox="1104 798 1161 840">512</div>

FIG. 5

RESOURCE GROUP TABLE (CLASS)	CLASS METHODS
ch_resource_group	A_sample_resource_group
failover_policy	cascading
failback_policy	autohoming
failback_action	update ch_resource_groups set next_node = \$_failback_node() where ch_resource_group = this ch_resource_group;
release_action	execute \$_resource_group_offline() ;
acquire_action	execute \$_resource_group_online() ;
current_node	
next_node	

FIG. 6